

ABSTRACT

A method is described for generating code for loading and performing other processing operations on a multi-dimensional data warehouse. Databases are defined as entity-relationship data models, using a modelling tool. A source file is created, containing instructions for processing the database, and including one or more high-level directives. At system build time, the source file is pre-processed, by replacing the directives with code, using information pulled from the data model, to generate a destination file containing the code for processing the database. At run time, the generated code is processed to replace run-time macros, indicating physical storage schemes. The generated code is then run, to process the database.